Chapter 7

Layer 3 VPN Configuration Guidelines

To configure Layer 3 virtual private network (VPN) functionality, you must enable VPN support on the provider edge (PE) router. You must also configure any provider (P) routers that service the VPN, and you must configure the customer edge (CE) routers so that their routes are distributed into the VPN.

To configure Layer 3 VPNs, you include statements at the [edit routing-instances] hierarchy level:

```
[edit]
routing-instances {
    routing-instance-name {
        description text;
        interface interface-name;
        instance-type vrf;
        route-distinguisher ( as-number:number | ip-address:number );
        vrf-import [ policy-names ];
        vrf-export [ policy-names ];
        vrf-table-label;
        protocols {
            bgp {
                bgp-configuration;
            }
            ospf {
                ospf-configuration;
            }
            pim {
                pim-configuration;
                vpn-group-address address;
            }
            rip {
                rip-configuration;
            }
        }
        routing-options {
            autonomous-system autonomous-system <loops number>;
            forwarding-table {
                export [ policy-names ];
            }
            interface-routes {
                rib-group group-name;
            }
            martians {
                destination-prefix match-type <allow>;
            }
            maximum-routes route-limit <log-only | threshold value>;
```

```
            options {
                syslog (level level | upto level);
            }
            rib routing-table {
                static {
                    defaults {
                        static-options;
                    }
                    route destination-prefix {
                        next-hop;
                        static-options;
                        }
                    }
                }
                martians {
                    destination-prefix match-type <allow>;
                }
                static {
                    defaults {
                        static-options;
                    }
                    route destination-prefix {
                        policy [ policy-names ];
                        static-options;
                    }
                }
            }
            router-id address;
            static {
                defaults {
                    static-options;
                }
                route destination-prefix {
                    policy [ policy-names ];
                    static-options;
                }
            }
        }
    }
}
```

For Layer 3 VPNs, only some of the statements in the [edit routing-instances] hierarchy are valid. For the full hierarchy, see the *JUNOS Internet Software Configuration Guide: Routing and Routing Protocols*.

In addition to these statements, you must enable a signaling protocol, internal Border Gateway Protocol (IBGP) sessions between the PE routers, and an interior gateway protocol (IGP) on the PE and provider routers.

By default, Layer 3 VPNs are disabled.

This chapter describes the following tasks for configuring VPNs:

For configuration examples, see “Layer 3 VPN Configuration Examples” on page 127.

Enable a Signaling Protocol

For Layer 3 VPNs to function, you must enable a signaling protocol on the PE routers. You can do one of the following:

Use LDP for VPN Signaling

To use Label Distribution Protocol (LDP) for VPN signaling, perform the following steps on the PE and provider routers:

1. Configure LDP on the interfaces in the core of the service provider’s network by including the ldp statement at the [edit protocols] hierarchy level. You need to configure LDP only on the interfaces between PE routers or between PE and provider routers. You can think of these as the “core-facing” interfaces. You do not need to configure LDP on the interface between the PE and CE routers.

```
[edit]
protocols {
    ldp {
        interface interface-name;
    }
}
```

2. Configure the Multiprotocol Label Switching (MPLS) address family on the interfaces on which you enabled LDP (the interfaces you configured in Step 1):

```
[edit]
interfaces {
    interface-name {
        unit logical-unit-number {
            family mpls;
        }
    }
}
```

Specify the interface name in the format *type-fpc/pic/port*.

3. Configure Open Shortest Path First (OSPF) or Intermediate System-to-Intermediate System (IS-IS) on each PE and provider router. You configure these protocols at the master instance of the routing protocol, not within the routing instance used for the VPN.

To configure OSPF, include the ospf statement at the [edit protocols] hierarchy level. At a minimum, you must configure a backbone area on at least one of the router's interfaces.

```
[edit]
protocols {
    ospf {
        area 0.0.0.0 {
            interface interface-name;
        }
    }
}
```

To configure IS-IS, include the isis statement at the [edit protocols] hierarchy level and configure the loopback interface and International Organization for Standardization (ISO) family at the [edit interfaces] hierarchy level. At a minimum, you must enable IS-IS on the router, configure a network entity title (NET) on one of the router's interfaces (preferably the loopback interface, lo0), and configure the ISO family on all interfaces on which you want IS-IS to run. When you enable IS-IS, Level 1 and Level 2 are enabled by default. The following is the minimum IS-IS configuration. In the address statement, *address* is the NET.

```
[edit]
interfaces {
    lo0 {
        unit logical-unit-number {
            family iso {
                address address;
            }
        }
    }
    type-fpc/pic/port {
        unit logical-unit-number {
            family iso;
        }
    }
}
protocols {
    isis {
        interface all;
    }
}
```

For more information about configuring OSPF and IS-IS, see the *JUNOS Internet Software Configuration Guide: Routing and Routing Protocols*.

Use RSVP for VPN Signaling

To use the Resource Reservation Protocol (RSVP) for VPN signaling, perform the following steps:

1. On each PE router, configure traffic engineering. To do this, you must configure an IGP that supports traffic engineering (either IS-IS or OSPF) and enable traffic engineering support for that protocol.

   To enable OSPF traffic engineering support, include the traffic-engineering statement at the [edit protocols ospf] hierarchy level:

   ```
   [edit protocols ospf]
   traffic-engineering {
       no-topology;
       shortcuts;
   }
   ```

   For IS-IS, traffic engineering support is enabled by default.

2. On each PE and provider router, enable RSVP on the router interfaces that participate in the label-switched path (LSP). On the PE router, these are the interfaces that are the ingress and egress points to the LSP. On the provider router, these are the interfaces that connect the LSP between the PE routers. Do not enable RSVP on the interface between the PE and the CE routers, because this interface is not part of the LSP.

   To configure RSVP on the PE and provider routers, include the interface statement at the [edit protocols rsvp] hierarchy level. Include one interface statement for each interface on which you are enabling RSVP.

   ```
   [edit protocols]
   rsvp {
       interface interface-name;
       interface interface-name;
   }
   ```

3. On each PE router, configure an MPLS LSP to the PE router that is the LSP's egress point. To do this, include the label-switched-path and interface statements at the [edit protocols mpls] hierarchy level.

   ```
   [edit protocols]
   mpls {
       label-switched-path path-name {
           to ip-address;
       }
       interface interface-name;
   }
   ```

   In the to statement, specify the address of the LSP's egress point, which is an address on the remote PE router.

   In the interface statement, specify the name of the interface (both the physical and logical portions). Include one interface statement for the interface associated with the LSP.

When you configure the same interface at the [edit interfaces] hierarchy level, you must also configure family mpls and family inet when configuring the logical interface:

```
[edit interfaces]
interface-name {
    unit logical-unit-number {
        family inet;
        family mpls;
    }
}
```

4. On all provider routers that participate in the LSP, enable MPLS by including the interface statement at the [edit mpls] hierarchy level. Include one interface statement for each connection to the LSP.

```
[edit]
mpls {
    interface interface-name;
    interface interface-name;
}
```

5. Enable MPLS on the interface between the PE and CE routers by including the interface statement at the [edit mpls] hierarchy level. Doing this allows the PE router to assign an MPLS label to traffic entering the LSP or to remove the label from traffic exiting the LSP.

```
[edit]
mpls {
    interface interface-name;
}
```

For information about configuring MPLS, see the *JUNOS Internet Software Configuration Guide: MPLS Applications*.

Configure an IGP on PE and Provider Routers

To allow the PE and provider routers to exchange routing information, you must either configure an IGP on all the routers or you must configure static routes. You configure the IGP on the master instance of the routing protocol process (rpd) (that is, at the [edit protocols] hierarchy level), not within the routing instance used for the VPN (that is, not at the [edit routing-instances] hierarchy level).

When you configure the PE router, do not configure any summarization of the PE router's loopback addresses at the area boundary. Each PE router's loopback address should appear as a separate route.

For information about configuring routing protocols and static routes, see the *JUNOS Internet Software Configuration Guide: Routing and Routing Protocols*.

Configure an IBGP Session between PE Routers

You must configure an IBGP session between PE routers to allow the PE routers to exchange information about routes originating and terminating in the VPN. To do this, include the family inet-vpn statement when configuring IBGP:

```
[edit protocols]
bgp {
    group group-name {
        type internal;
        local-address ip-address;
        family inet-vpn {
            unicast;
        }
        neighbor ip-address;
    }
}
```

The family inet-vpn statement indicates that the IBGP session is for the VPN.

The IP address in the local-address statement is the address of the loopback interface (lo0) on the local PE router. The IBGP session for VPNs runs through the loopback address. (You must also configure the lo0 interface at the [edit interfaces] hierarchy level.)

The IP address in the neighbor statement is the loopback address of the neighboring PE router. If you are using RSVP signaling, this IP address is the same address you specify in the to statement at the [edit mpls label-switched-path *lsp-path-name*] hierarchy level when you configure the MPLS LSP.

Configure Routing Instances for Layer 3 VPNs on PE Routers

To configure routing instances for Layer 3 VPNs, include the routing-instances statement at the [edit] hierarchy level. You configure VPN routing instances only on PE routers.

```
[edit]
routing-instances {
    routing-instance-name {
        description text;
        interface interface-name;
        instance-type vrf;
        route-distinguisher ( as-number:number | ip-address:number );
        vrf-import [ policy-name ];
        vrf-export [ policy-name ];
        vrf-target {
            export community-name;
            import community-name;
        }
    }
}
```

For the VPN to function, you must include the instance-type, interface, route-distinguisher, vrf-import, and vrf-export statements in the routing instance configuration on the PE router. The vrf-table-label statement is optional.

The following sections describe how to configure VPN routing instances:

Configure the Description

To provide a textual description for the routing instance, include the description statement at the [edit routing-instances *routing-instance-name*] hierarchy level. Enclose any descriptive text that includes spaces in quotation marks (" "). Any descriptive text you include is displayed in the output of the show route instance detail command and has no effect on the operation of the routing instance.

```
[edit routing-instances routing-instance-name]
description text;
```

Configure the Instance Type

Each PE router uses a VPN routing and forwarding (VRF) table for distributing routes within the VPN. To create the VRF table on the PE router, include the instance-type statement at the [edit routing-instances *routing-instance-name*] hierarchy level, specifying the instance type as vrf:

```
[edit routing-instances routing-instance-name]
instance-type vrf;
```

Configure Interfaces for VPN Routing

On each PE router, you must configure an interface over which the VPN traffic travels between the PE and CE routers. To do this, include the interface statement at the [edit routing-instances *routing-instance-name*] hierarchy level:

```
[edit routing-instances routing-instance-name]
interface interface-name;
```

Specify both the physical and logical portions of the interface name, in the following format:

```
physical.logical
```

For example, in so-1/2/1.0, so-1/2/1 is the physical portion of the interface name and 0 is the logical portion. If you do not specify the logical portion of the interface name, 0 is used.

A logical interface can be associated with only one routing instance.

When you configure the same interface at the [edit interfaces] hierarchy level, you must also configure family inet when configuring the logical interface:

```
[edit interfaces]
interface-name {
    unit logical-unit-number {
        family inet;
    }
}
```

If you downgrade JUNOS to version 5.5 or earlier, the VPN configuration might become invalid. VPN interfaces between PE and CE routers formerly required the family mpls statement to be configured.

When you configure carrier-of-carriers VPNs, you need to configure the family mpls statement in addition to the family inet statement for the interfaces between the PE and CE routers. For carrier-of-carriers VPNs, configure the logical interface as follows:

```
[edit interfaces]
interface-name {
    unit logical-unit-number {
        family inet;
        family mpls;
    }
}
```

If you configure family mpls on the logical interface and then configure this interface for a non-carrier-of-carriers vrf routing instance, the family mpls statement is automatically removed from the configuration for the logical interface, since it is not needed.

If you enable a routing protocol on all instances by specifying interfaces all when configuring the master instance of the protocol at the [edit protocols] hierarchy level, and if you configure a specific interface for VPN routing at the [edit routing-instances *routing-instance-name*] hierarchy level, the latter interface statement takes precedence and the interface is used exclusively for the VPN.

If you explicitly configure the same interface name at both the [edit protocols] and [edit routing-instances *routing-instance-name*] hierarchy levels, when you try to commit the configuration, it will fail.

Configure a Logical Unit on the Loopback Interface

You can configure a logical unit on the loopback interface into each VRF routing instance you have configured on the router. This is only possible on Layer 3 VPNs (VRF routing instances). Associating a VRF routing instance with a logical unit on the loopback interface allows you to easily identify the VRF. This is useful for troubleshooting, allowing you to ping a remote CE router from a local PE router in a Layer 3 VPN. See "Ping a Remote CE Router from a PE Router" on page 123 for more information.

You can also configure a firewall filter for the logical unit on the loopback interface, allowing you to filter traffic for the VRF routing instance associated with it.

The following describes how firewall filters affect the VRF routing instance depending on whether they are configured on the default loopback interface, the VRF routing instance, or some combination of the two. Note that "default loopback interface" refers to lo0.0 (associated with the default routing table) and "VRF loopback interface" refers to lo0.*n* which is configured in the VRF routing instance.

If you configure Filter A on the default loopback interface and Filter B on the VRF loopback interface, the VRF routing instance uses Filter B.

If you configure Filter A on the default loopback interface, but do not configure a filter on the VRF loopback interface, the VRF routing instance does not use a filter.

If you configure Filter A on the default loopback interface but do not even configure a VRF loopback interface, the VRF routing instance uses Filter A.

To configure a logical unit on the loopback interface, configure the unit statement at the [edit interfaces lo0] hierarchy level:

```
[edit interfaces]
lo0 {
    unit number {
        family inet {
            address address;
        }
}
```

To associate a firewall filter with the logical unit on the loopback interface, include the following statements at the [edit interfaces lo0 unit *unit-number* family inet] hierarchy level:

```
[edit interfaces lo0 unit unit-number family inet]
filter {
    input filter-name;
}
```

You also need to include the lo0.*n* interface in the configuration for the VRF routing instance at the [edit routing-instances *routing-instance-name*] hierarchy level:

```
[edit]
routing-instances {
    routing-instance-name {
        interface lo0.n;
    }
}
```

For more information on how to configure firewall filters, see the *JUNOS Internet Software Configuration Guide: Policy Framework*.

Configure the Route Distinguisher

The configuration for a route distinguisher in a Layer 3 VPN is identical to the configuration in a Layer 2 VPN. See “Configure the Route Distinguisher” on page 21.

Configure Policy for the PE Router’s VRF Table

On each PE router, you must define policies that define how routes are imported into and exported from the router’s VRF table. In these policies, you must define the route target and you can optionally define the route origin.

The following sections describe how to configure policy for the VRF tables:

Configure the Route Target

In the import and export policies for the PE router’s VRF table, you must define the route target, which defines which VPN the route is part of. To do this, include the target option in the community statement at the [edit policy-options] hierarchy level:

```
[edit policy-options]
community name members target:community-id;
```

name is the name of the community.

community-id is the identifier of the community. You specify it in one of the following formats:

as-number:*number*, where *as-number* is an autonomous system (AS) number (a 2-byte value) and *number* is a 4-byte community identifier. The AS number can be in the range 1 through 65,535. We recommend that you use an Internet Assigned Numbers Authority (IANA)-assigned, nonprivate AS number, preferably the Internet service provider’s (ISP’s) own or the customer’s own AS number. The community identifier can be a number in the range 0 through $2^{32} – 1$.

ip-address:*number*, where *ip-address* is an Internet Protocol Version 4 (IPv4) address (a 4-byte value) and *number* is a 2-byte community identifier. The IP address can be any globally unique unicast address. We recommend that you use the address that you configure in the router-id statement, which is a nonprivate address in your assigned prefix range. The community identifier can be a number in the range 1 through 65,535.

Configure the Route Origin

In the import and export policies for the PE router's VRF table, you can optionally define the route origin (also known as the site of origin), which identifies the set of routes learned from a particular CE site. This attribute ensures that a route learned from a particular site through a particular PE-CE connection is not distributed back to the site through a different PE-CE connection. It is particularly useful if you are using the Border Gateway Protocol (BGP) as the routing protocol between the PE and CE routers and if different sites in the VPN have been assigned the same AS numbers.

To configure a route origin, complete the following steps:

1. Include the origin option in the community statement at the [edit policy-options] hierarchy level:

   ```
   [edit policy-options]
   community name members origin:community-id;
   ```

   name is the name of the community.

   community-id is the identifier of the community. You specify it in one of the following format:

   as-number:*number*, where *as-number* is an AS number (a 2-byte value) and *number* is a 4-byte community identifier. The AS number can be in the range 1 through 65,535. We recommend that you use an IANA-assigned, nonprivate AS number, preferably the ISP's own or the customer's own AS number. The community identifier can be a number in the range 0 through $2^{32} - 1$.

   ip-address:*number*, where *ip-address* is an IPv4 address (a 4-byte value) and *number* is a 2-byte community identifier. The IP address can be any globally unique unicast address. We recommend that you use the address that you configure in the router-id statement, which is a nonprivate address in your assigned prefix range. The community identifier can be a number in the range 1 through 65,535.

2. Include the community in the import policy for the PE router's VRF table by configuring the community statement with the *community-id* defined in Step 1 at the [edit policy-options policy-statement *import-policy-name* term *import-term-name* from] hierarchy level. See "Configure Import Policy for the PE Router's VRF Table" on page 90.

3. Include the community in the export policy for the PE router's VRF table by configuring the community statement with the *community-id* defined in Step 1 at the [edit policy-options policy-statement *export-policy-name* term *export-term-name* then] hierarchy level. See "Configure Export Policy for the PE Router's VRF Table" on page 91.

Configure Import Policy for the PE Router's VRF Table

Each VPN must have a policy that defines how routes are imported into the PE router's VRF table. An import policy is applied to routes received from other PE routers in the VPN. A policy must evaluate all routes received over the IBGP session with the peer PE router. If the routes match the conditions, the route is installed in the PE router's *routing-instance-name*.inet.0 VRF table. An import policy must contain a second term that rejects all other routes.

Unless an import policy contains only a then reject statement, it must include a reference to a community. Otherwise, when you try to commit the configuration, the commit fails. Note that you can configure multiple import policies.

An import policy determines what to import to a specified VRF table based on the VPN routes learned from the remote PE routers through IBGP. The IBGP session is configured at the [edit protocols bgp] hierarchy level. If you also configure an import policy at the [edit protocols bgp] hierarchy level, the import policies at the [edit policy-options] hierarchy level and the [edit protocols bgp] hierarchy level are combined through a logical AND operation. This allows you to filter traffic as a group.

To configure an import policy for the PE router's VRF table, follow these steps:

1. To define an import policy, include the policy-statement statement at the [edit policy-options] hierarchy level. For all PE routers, an import policy must always include the following, at a minimum:

```
[edit]
policy-options {
    policy-statement import-policy-name {
        term import-term-name {
            from {
                protocol bgp;
                community community-id;
            }
            then accept;
        }
        term term-name {
            then reject;
        }
    }
}
```

   The *import-policy-name* policy evaluates all routes received over the IBGP session with the other PE router. If the routes match the conditions in the from statement, the route is installed in the PE router's *routing-instance-name*.inet.0 VRF table. The second term in the policy rejects all other routes.

   For more information about creating policies, see the *JUNOS Internet Software Configuration Guide: Routing and Routing Protocols*.

2. To configure an import policy, include the vrf-import statement at the [edit routing-instances *routing-instance-name*] hierarchy level:

```
[edit routing-instances routing-instance-name]
vrf-import [ import-policy-name ];
```

Configure Export Policy for the PE Router's VRF Table

Each VPN must have a policy that defines how routes are exported from the PE router's VRF table. An export policy is applied to routes sent to other PE routers in the VPN. An export policy must evaluate all routes received over the routing protocol session with the CE router. (This session can use the BGP, OSPF, or Routing Information Protocol (RIP) routing protocols, or static routes.) If the routes match the conditions, the specified community target (which is the route target) is added to them and they are exported to the remote PE routers. An export policy must contain a second term that rejects all other routes.

Export policies defined within the VPN routing instance are the only export policies that apply to the VRF table. Any export policy that you define on the IBGP session between the PE routers has no effect on the VRF table. Note that you can configure multiple export policies.

To configure an export policy for the PE router's VRF table, follow these steps:

1. To define an export policy, include the policy-statement statement at the [edit policy-options] hierarchy level. For all PE routers, an export policy must distribute VPN routes to and from the connected CE routers in accordance with the type of routing protocol that you configure between the CE and PE routers within the routing instance. An export policy must always include the following, at a minimum:

```
[edit]
policy-options {
    policy-statement export-policy-name {
        term export-term-name {
            from protocol (bgp | ospf | rip | static);
            then {
                community add community-id;
                accept;
            }
        }
        term term-name {
            then reject;
        }
    }
}
```

   The *export-policy-name* policy evaluates all routes received over the routing protocol session with the CE router. (This session can use the BGP, OSPF, or RIP routing protocol, or static routes.) If the routes match the conditions in the from statement, the community target specified in the then community add statement is added to them and they are exported to the remote PE routers. The second term in the policy rejects all other routes.

   For more information about configuring routing within the routing instance, see "Configure VPN Routing between the PE and CE Routers" on page 96. For more information about creating policies, see the *JUNOS Internet Software Configuration Guide: Routing and Routing Protocols*.

2. To apply the policy, include the vrf-export statement at the [edit routing-instances *routing-instance-name*] hierarchy level:

```
[edit routing-instances routing-instance-name]
vrf-export export-policy-name;
```

Apply Both the VRF Export and the BGP Export Policies

When you apply a VRF export policy as described in "Configure Export Policy for the PE Router's VRF Table" on page 91, routes from Layer 2 VPN (l2vpn) or Layer 3 VPN (vrf) routing instances are advertised to other PE routers based on this policy, while the BGP export policy is ignored.

If you configure the vpn-apply-export statement, both the VRF export and BGP group or neighbor export policies are applied (VRF first, then BGP) before routes are advertised in the vrf or l2vpn routing tables to other PE routers.

Configure the vpn-apply-export statement at the [edit protocols bgp], [edit protocols bgp group *group-name*], or [edit protocols bgp group *group-name* neighbor *neighbor*] hierarchy level:

```
[edit]
vpn-apply-export;
```

Configure a VRF Target

Before JUNOS 5.5, you needed to configure VRF import and export policies for each VPN routing instance on a PE router. These policies control redistribution of routes between the VRF table and BGP.

In the current JUNOS release, the vrf-target statement simplifies this configuration. Configuring a VRF target community using the vrf-target statement causes default VRF import and export policies to be generated which, respectively, accept and tag routes with the specified target community. You can still create more complex policies by explicitly configuring VRF import and export policies. Note that these policies would override the default policies generated when you configure the vrf-target statement.

If you do not configure the import and export options of the vrf-target statement, the specified community string is applied in both directions. The import and export keywords give you more flexibility, allowing you to specify a different community for each direction.

An example of how you might configure the vrf-target statement follows:

```
[edit routing-instances sample]
vrf-target target:69:102;
```

Note that the syntax for the VRF target community is not a name. You must specify it in the format target:*x*:*y*. A community name cannot be specified because this would also require you to configure the community members for that community using the policy-options statement. If you define the policy-options statements, then you can just configure VRF import and export policies as usual. The purpose of the vrf-target statement is to simplify the configuration by allowing you to configure most statements at the [edit routing-instances] hierarchy level.

To configure a VRF target, include the vrf-target statement at the [edit routing-instances *routing-instance-name*] hierarchy level:

```
[edit]
routing-instances {
    routing-instance-name {
        vrf-target community;
    }
}
```

To configure the vrf-target statement with the export and import options, include the following statements at the [edit routing-instances *routing-instance-name*] hierarchy level

```
[edit]
routing-instances {
    routing-instance-name {
        vrf-target {
            export community-name;
            import community-name;
        }
    }
}
```

Filter Traffic Based on the IP Header

The vrf-table-label statement makes it possible to map the inner label to a specific VRF and thus allow the examination of the encapsulated IP header at an egress VPN router. You might want to enable this functionality so you can do either of the following:

Forward traffic on a PE-router-to-CE-device interface, in a shared medium, where the CE device is a Layer 2 switch without IP capabilities (for example, a metro Ethernet switch).

The first lookup is done on the VPN label to determine which VRF table to refer to, and the second lookup is done on the IP header to determine how to forward packets to the correct end hosts on the shared medium.

Perform egress filtering at the egress PE router.

The first lookup on the VPN label is done to determine which VRF table to refer to, and the second lookup is done on the IP header to determine how to filter and forward packets. You can enable this functionality by configuring output filters on the VRF interfaces.

The vrf-table-label statement is supported on the M-series platforms only. It is not supported on the T-series platforms.

When you use the vrf-table-label statement to configure a VRF table, a label-switched interface (LSI) logical interface label is created and mapped to the VRF. You perform this configuration at the [edit routing-instances *routing-instance-name*] hierarchy level.

Any routes configured in a VRF with the vrf-table-label statement are advertised with the LSI logical interface label allocated for the VRF. When packets for this VPN arrive on a core-facing interface, they are treated as if the enclosed IP packet arrived on the LSI interface and are then forwarded and filtered based on the correct table.

To filter traffic based on the IP header, include the vrf-table-label statement at the [edit routing-instances *routing-instance-name*] hierarchy level:

```
[edit routing-instances routing-instance-name]
vrf-table-label;
```

Do not use the vrf-table-label statement for source class usage/destination class usage (SCU/DCU) configurations. For information on SCU/DCU configuration, see the *JUNOS Internet Softw are Configur ation Guide: Interf aces and Class of Service* .

Egress Filtering Options and Limit ations

Egress filtering (which allows egress Layer 3 VPN PE routers to perform lookups on the VPN label and IP header at the same time) can be enabled by including the vrf-table-label statement at the [edit routing-instances *instance-name*] hierarchy level. However, this feature works only for non-channelized Point-to-Point Protocol/High-level Data Link Control (PPP/HDLC) core-facing SONET interfaces. Note that the vrf-table-label statement cannot be configured for the 10-port E1 Physical Interface Card (PIC) or for aggregated interfaces. There is no restriction on CE-router-to-PE-router interfaces.

You can also enable egress filtering by configuring a VPN tunnel (VT) interface on routers equipped with a Tunnel Services PIC. When you enable egress filtering this way, there is no restriction on the type of core-facing interface used. There is also no restriction on the type of CE-router-to-PE-router interface used.

You cannot configure a VT interface and the vrf-table-label statement at the same time.

Configure a VPN Tunnel for VRF Table Lookup

You can configure a VPN tunnel to facilitate VRF table lookup based on MPLS labels. You might want to enable this functionality to forward traffic on a PE-router-to-CE-device interface in a shared medium, where the CE device is a Layer 2 switch without IP capabilities (for example, a metro Ethernet switch), or to perform egress filtering at the egress PE router.

For more information on VPN tunnels and VT interfaces, see the *JUNOS Internet Softw are Configur ation Guide: Interf aces and Class of Service* .

Configure VPN Routing between the PE and CE Routers

For the PE router to distribute VPN-related routes to and from connected CE routers, you must configure routing within the VPN routing instance. You can configure a routing protocol—BGP, OSPF, or RIP—or you can configure static routing. For the connection to each CE router, you can configure only one type of routing.

This section describes how to do the following tasks:

Configure BGP between the PE and CE Routers

To configure BGP as the routing protocol between the PE and the CE routers, include the bgp statement at the [edit routing-instances *routing-instance-name* protocols] hierarchy level:

```
[edit routing-instances routing-instance-name protocols]
bgp {
    group group-name {
        peer-as as-number;
        neighbor ip-address;
    }
}
```

Configure OSPF between the PE and CE Routers

You can configure OSPF to distribute VPN-related routes between PE and CE routers.

To configure OSPF as the routing protocol between a PE and CE router, include the protocols ospf statement at the [edit routing-instances *routing-instance-name* protocols] hierarchy level:

```
[edit routing-instances routing-instance-name protocols]
ospf {
    area area {
        interface interface-name;
    }
}
```

Configure an OSPF Domain ID

For most OSPF configurations involving Layer 3 VPNs, you do not need to configure an OSPF domain ID. However, for a Layer 3 VPN connecting multiple OSPF domains, configuring OSPF domain IDs can help you control link-state advertisement (LSA) translation (for Type 3 and Type 5 LSAs) between the OSPF domains and back-door paths. The default OSPF domain ID is 0.0.0.0. Each VRF table in a PE router associated with an OSPF instance is configured with the same OSPF domain ID.

When a PE router receives a route, it redistributes and advertises the route either as a Type 3 LSA or as a Type 5 LSA, depending on the following:

If the receiving PE router sees a Type 3 route with a matching domain ID, the route is redistributed and advertised as a Type 3 LSA.

If the receiving PE router sees a Type 3 route without a domain ID (the extended attribute field of the route's BGP update does not include a domain ID), the route is redistributed and advertised as a Type 3 LSA.

If the receiving PE router sees a Type 3 route with a non-matching domain ID, the route is redistributed and advertised as a Type 5 LSA.

If the receiving PE router sees a Type 3 route with a domain ID but the receiving PE router does not have a domain ID configured, the route is redistributed and advertised as a Type 5 LSA.

If the receiving PE router sees a Type 5 route, the route is redistributed and advertised as a Type 5 LSA, irrespective of the domain ID.

To configure an OSPF domain ID, include the domain-id statement at the [edit routing-instances *routing-instance-name* protocols ospf] hierarchy level:

```
[edit routing-instances routing-instance-name protocols]
ospf {
    domain-id domain ID;
}
```

You can set a VPN tag for the OSPF external routes generated by the PE router. This is used to prevent looping when a domain ID is used as an alternate route preference. By default, this tag is automatically calculated and needs no configuration. To configure the domain VPN tag for Type 5 LSAs, include the domain-vpn-tag *number* statement at the [edit routing-instances *routing-instance-name* protocols ospf] hierarchy level:

```
[edit routing-instances routing-instance-name protocols]
ospf {
    domain-vpn-tag number;
}
```

The range is 1 through 4,294,967,295. If you set VPN tags manually, you must set the same value for all PE routers in the VPN.

Hub-and-Spoke Layer 3 VPNs and OSPF Domain ID

The default behavior of an OSPF domain ID can cause the following problems for hub-and-spoke Layer 3 VPNs using OSPF between the PE and CE routers:

PE routers set the down (DN) bit on all OSPF summary LSAs originating from area 0. PE routers are designated as area 0 by default because of the OSPF domain ID. When a PE router receives a summary LSA with the DN bit set, it drops the LSA. This is done to prevent routing loops.

For a hub-and-spoke Layer 3 VPN, when the hub PE router generates an OSPF summary LSA, it also sets the DN bit before sending it to the hub CE router. When the hub CE router sends the LSA back to the PE router, the PE router drops the LSA because the DN bit is set. Routes aggregated within the CE router are not affected by this problem.

PE routers generating external LSAs learned from BGP updates set the vpn-route-tag field to a value derived from the PE router's AS number and an arbitrary tag. When a PE router receives an external LSA with a vpn-route-tag field that matches its own vpn-route-tag field, it drops the LSA. This is done to prevent routing loops.

For a hub-and-spoke Layer 3 VPN, an external LSA originated by a hub PE router is sent to the hub CE router, which then sends it back to the same PE router. Because the vpn-route-tag field matches the PE router's vpn-route-tag field, the LSA is dropped. Routes aggregated within the CE router are not affected by this problem.

For hub-and-spoke Layer 3 VPNs using OSPF between the PE and CE routers to work, you need to configure the following on the hub PE router:

Configure the disable statement at the [edit routing-instances *routing-instance-name* protocols ospf] hierarchy level on the routing instance for the hub CE router. This removes area 0 from the PE router, allowing the PE router to forward LSAs without setting the DN bit. When an LSA comes back from the hub CE router, the PE router can install it because the DN bit is not set.

Configure 0 for the vpn-route-tag statement at the [edit routing-instances *routing-instance-name* protocols ospf] hierarchy level on the routing instance for the spoke CE router. This removes any VPN route tags that are set on the external LSAs, preventing a VPN route tag match and allowing the PE router to install the LSA.

Compatibility with JUNOS Releases before 5.3

For JUNOS release 5.3, the format for domain-id, an extended community type defined in the BGP extended community attribute field, was modified to comply with the Internet Engineering Task Force (IETF) draft draft-rosen-vpns-ospf-bgp-mpls (available at http://www.ietf.org/). JUNOS releases prior to 5.3 continue to use the previously supported vendor-specific formats.

The OSPF domain ID format is incompatible between JUNOS 5.3 or later and JUNOS 5.2 or earlier. For OSPF domain IDs to function properly between a PE router running JUNOS 5.3 or later and a PE router running JUNOS 5.2 or earlier, you need to define the extended community type for the BGP extended community attribute field as domain-id-vendor (instead of as domain-id). This is part of the policy-options configuration for the OSPF domain ID configured at the [edit policy-options community vrf_export_attributes members] hierarchy level:

```
[edit policy-options community vrf_export_attributes members]
domain-id-vendor:ip-address
```

You also need to configure the route-type-community statement with the vendor option at the [edit routing-instances *routing-instance-name* protocols ospf] hierarchy level:

```
[edit routing-instances routing-instance-name protocols]
ospf {
    route-type-community vendor;
}
```

The default value for the route-type-community statement is iana.

Example Configur ations for Compatibility with JUNOS R eleases bef ore 5.3

The following example shows a configuration of the policy options for a PE router. The PE router has an OSPF domain ID configured.

It needs to be compatible with a router running a pre-5.3 version of JUNOS software. As a part of the community statement configuration, specify domain-id-vendor for the attribute that assigns the domain ID instead of domain-id:

```
[edit]
policy-options {
    policy-statement vrf_import_routes {
        term a {
            from {
                protocol bgp;
                community vrf_import_attributes;
            }
            then accept;
        }
        term b {
            then reject;
        }
    }
    policy-statement vrf_export_routes {
        term a {
            from protocol ospf;
            then {
                community add vrf_export_attributes;
                accept;
            }
        }
        term b {
            then reject;
        }
    }
    community vrf_export_attributes members [ target:10.19.2.0:5 domain-id-vendor:1.2.3.4:0 ];
    community vrf_import_attributes members target:10.19.1.0:5;
}
```

The following example shows a configuration for a routing instance on a PE router. The PE router has an OSPF domain ID configured. It needs to be compatible with a router running an earlier version of JUNOS software. The configuration includes the route-type-community statement with the vendor option. This is so the PE router receiving the route knows how to parse the incoming BGP attribute field containing the domain ID.

The example configuration follows:

```
[edit]
routing-instances {
    CE_A {
        instance-type vrf;
        interface fe-1/0/0.0;
        route-distinguisher 10.255.25.270:1;
        vrf-import vrf_import_routes;
        vrf-export vrf_export_routes;
        protocols {
            ospf {
                route-type-community vendor;
                domain-id 1.2.3.4;
                export vrf_import_routes;
                area 0.0.0.0 {
                    interface fe-1/0/0.0;
                }
            }
        }
    }
}
```

Configure RIP between the PE and CE Routers

For a Layer 3 VPN, you can configure RIP on the PE router to learn the routes of the CE router or to propagate the routes of the PE router to the CE router. RIP routes learned from neighbors configured at any [edit routing-instances] hierarchy level are added to the routing instance's inet table (*instance_name*.inet.0).

To configure RIP as the routing protocol between the PE and the CE router, include the rip statement at the [edit routing-instances *routing-instance-name* protocols] hierarchy level:

```
[edit routing-instances routing-instance-name protocols]
rip {
    group group-name {
        neighbor interface-name;
    }
}
```

To install routes learned from a RIP routing instance to multiple routing tables, configure the rib-group statement at the [edit protocols rip] hierarchy level or at the [edit routing-instances *routing-instance-name* protocols rip] hierarchy level:

```
[edit protocols rip]
rib-group inet group-name;
group group-name {
    neighbor interface-name;
}
```

To configure a routing table group, configure the rib-group statement at the [edit routing-options] hierarchy level.

To add a routing table to a routing table group, you need to configure the the import-rib statement at the [edit routing-options rib-groups *group-name*] hierarchy level. The first routing table name specified under the import-rib statement must be the name of the routing table you are configuring. See the *JUNOS Internet Software Configuration Guide: Routing and Routing Protocols* for more information about how to configure routing tables and routing table groups.

Configure the import-rib statement at the [edit routing-options rib-groups *group-name*] hierarchy level as follows:

```
[edit routing-options rib-groups group-name]
import-rib [group-name]
```

Configure Static Routes between the PE and CE Routers

To configure a static route between the PE and the CE routers, include the routing-options static statement at the [edit routing-instances *routing-instance-name* routing-options] hierarchy level:

```
[edit routing-instances routing-instance-name routing-options]
static {
    route destination-prefix {
        next-hop;
        static-options;
    }
}
```

For more information about configuring routing protocols and static routes, see the *JUNOS Internet Software Configuration Guide: Routing and Routing Protocols*.

Limit the Routes Accepted from a CE Router

A route limit sets an upper limit for the number of prefixes installed into routing tables. You can use route limits to curtail the number of routes received from a CE router in a VPN. A route limit applies only to dynamic routing protocols, and is not applicable to static or interface routes.

To limit the number of routes accepted by a PE router from a CE router, include the maximum-routes statement at the [edit routing-instances *routing-instance-name* routing-options] hierarchy level:

```
[edit routing-instances routing-instance-name routing-options]
maximum-routes route-limit <log-only | threshold value>;
```

There are two modes for route limits: advisory (set with the log-only option) and mandatory. An advisory limit triggers only warnings. The log messages are rate-limited to once every 30 seconds. A mandatory limit, in addition to triggering a warning message, rejects any additional routes after the threshold is reached. The threshold value is a percentage of the route limit at which warning messages are logged.

Setting a route limit might result in unpredictable dynamic routing protocol behavior.

Configure IPv6 between the PE and CE Routers

You can configure IPv6 between the PE and CE routers of a Layer 3 VPN. The PE router must have the PE router to PE router BGP session configured with the family inet6-vpn statement. The CE router must be capable of receiving IPv6 traffic. You can configure BGP or static routes between the PE and CE routers.

To configure IPv6 VPNS between the PE routers, complete the following steps:

Configure IPv6 on the PE Router on page 102

Configure BGP or Static Routes on the PE Router on page 102

Configure IPv6 on the Interfaces on page 104

Configure IPv6 on the PE Router

To configure IPv6 between the PE and CE routers, include the following statements at the [edit protocols bgp group *group-name*] hierarchy level on the PE router:

```
[edit protocols bgp group group-name]
family inet6-vpn {
    (unicast | multicast | any) {
        prefix-limit maximum prefix-limit;
        rib-group rib-group-name;
    }
}
```

Configure the ipv6-tunneling statement at the [edit protocols mpls] hierarchy level:

```
[edit protocols mpls]
ipv6-tunneling;
```

Configure BGP or Static Routes on the PE Router

You must configure either BGP or static routes for the connection between the PE and CE routers in the Layer 3 VPN. You can configure BGP to handle just IPv6 routes or both IPv4 and IPv6 routes.

For more information about IPv6, see the *JUNOS Internet Software Configuration Guide: Routing and Routing Protocols*.

The following sections describe how to configure BGP and static routes:

Configure BGP on the PE Router to Handle IPv6 Routes on page 103

Configure BGP on the PE Router to Handle IPv4 and IPv6 Routes on page 103

Configure Static Routes on the PE Router on page 104

Configure BGP on the PE Router to Handle IPv6 Routes

Configure BGP in the Layer 3 VPN routing instance to handle IPv6 routes at the [edit routing-instances *routing-instance-name* protocols bgp] hierarchy level:

```
[edit]
routing-instances routing-instance-name {
    protocols {
        bgp {
            group group-name {
                local-address IPv6-address;
                family inet6 {
                    unicast;
                }
                peer-as as-number;
                neighbor IPv6-address;
            }
        }
    }
}
```

Configure BGP on the PE Router to Handle IPv4 and IPv6 Routes

Configure BGP in the Layer 3 VPN routing instance to handle both IPv4 and IPv6 routes at the [edit routing-instances *routing-instance-name* protocols bgp] hierarchy level:

```
[edit]
routing-instances routing-instance-name {
    protocols {
        bgp {
            group group-name {
                local-address IPv4-address;
                family inet {
                    unicast;
                }
                family inet6 {
                    unicast;
                }
                peer-as as-number;
                neighbor address;
            }
        }
    }
}
```

Configure Static Routes on the PE Router

Configure a static route to the CE router in the Layer 3 VPN routing instance at the [edit routing-instances *routing-instance-name* routing-options rib *routing-table-group-name*] hierarchy level:

```
[edit routing-instances routing-instance-name]
routing-options {
    rib routing-table-group-name.inet6.0 {
        static {
            defaults {
                static-options;
            }
        }
    }
}
```

Configure IPv6 on the Interfaces

You need to configure IPv6 on the PE router interfaces to the CE routers and on the CE router interfaces to the PE routers.

To configure the interface to handle IPv6 routes, include the family inet6 statement under the [edit interfaces *interface-name* unit *unit-number*] hierarchy level:

```
[edit]
interfaces {
    interface-name {
        unit unit-number {
            family inet6 {
                address IPv6-address;
            }
        }
    }
}
```

If you have configured the Layer 3 VPN to handle both IPv4 and IPv6 routes, you need to configure the interface to handle both IPv4 and IPv6 routes:

```
[edit]
interfaces {
    interface-name {
        unit unit-number {
            family inet {
                address IPv4-address;
            }
            family inet6 {
                address IPv6-address;
            }
        }
    }
}
```

Configure EBGP or IBGP Multihop between PE and CE Routers

You can configure an external BGP (EBGP) or internal BGP (IBGP) multihop session between the PE and CE routers of a Layer 3 VPN. This allows you to have one or more routers between the PE and CE routers. Using IBGP between PE and CE routers does not require the configuration of any additional statements. However, using EBGP between the PE and CE routers requires the configuration of the multihop statement.

To configure an external BGP multihop session for the connection between the PE and CE routers, include the multihop statement at the [edit routing-instances *routing-instance-name* protocols bgp], [edit routing-instances *routing-instance-name* protocols bgp group *group-name*], or [edit routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor *address*] hierarchy level:

```
multihop <ttl-value>;
```

Configure VPN Graceful Restart

To enable VPN graceful restart, configure the graceful-restart statement at the [edit routing-options] hierarchy level on the PE router:

```
[edit routing-options]
graceful-restart {...}
```

Also include the graceful-restart statement at the [edit routing-instance *routing-instance-name* routing-options] hierarchy level on the PE router:

```
[edit routing-instance routing-instance-name routing-options]
graceful-restart {...}
```

Configure Multicast over Layer 3 VPNs

You can configure a Layer 3 VPN to support multicast traffic using the Protocol Independent Multicast (PIM) routing protocol. To support multicast, you need to configure PIM on routers within the VPN and within the service provider's network.

Each PE router configured to run multicast over Layer 3 VPNs must have a Tunnel PIC. A Tunnel PIC is also required on the provider routers that act as rendezvous points (RPs). Tunnel PICs are also needed on all the CE routers acting as designated routers (first-hop/last-hop routers) or as RPs, just as they are in non-VPN PIM environments.

Configure the master PIM instance at the [edit protocols pim] hierarchy level on the CE and PE routers. You also need to configure a PIM instance for the Layer 3 VPN at the [edit routing-instances *routing-instance-name* protocols pim] hierarchy level on the PE router. This creates a PIM instance for the indicated routing instance.

For information about how to configure PIM, see the *JUNOS Internet Softw are Configur ation Guide: Multicast* .

The vpn-group-address statement is unique to a Layer 3 VPN PIM configuration. You use this statement to configure the group address for the VPN in the service provider's network. This address should be unique for each VPN. It ensures that multicast traffic is transmitted only to the specified VPN.

Configure the vpn-group-address statement at the [edit routing-instances *routing-instance-name* protocols pim] hierarchy level:

```
[edit routing-instances routing-instance-name protocols]
pim {
    vpn-group-address address;
}
```

The rest of the Layer 3 VPN configuration for multicast is conventional and is described in other sections of this manual. Most of the specific configuration tasks needed to activate multicast in a VPN environment involve PIM. For more information about how to configure PIM and multicast in JUNOS, including an example of how to configure multicast over Layer 3 VPNs, see the *JUNOS Internet Software Configuration Guide: Multicast*.

Configure a GRE Tunnel Interface for Layer 3 VPNs

JUNOS software allows you to configure a generic routing encapsulation (GRE) tunnel between the PE and CE routers for a Layer 3 VPN. The GRE tunnel can have one or more hops.

For more information about how to configure tunnel interfaces, see the *JUNOS Internet Software Configuration Guide: Interfaces and Class of Service*.

To configure a GRE tunnel between the PE and CE routers for a Layer 3 VPN, complete the procedures in the following sections:

Configure the GRE Tunnel Interface on the PE Router on page 106

Configure the GRE Tunnel Interface on the CE Router on page 107

Configure the GRE Tunnel Interface on the PE Router

Configure the GRE tunnel interface on the PE router as follows:

```
[edit]
interfaces {
    interface-name {
        unit 0 {
            tunnel {
                source address;
                destination address;
            }
            family inet {
                address address;
            }
        }
    }
}
```

By default, the tunnel destination address is assumed to be in the default Internet routing table, inet.0. If the tunnel destination address is not in inet.0, you need to specify which routing table to search for the tunnel destination address by configuring the routing-instance statement. This is the case if the tunnel encapsulating interface is also configured under the routing instance.

Configure the GRE tunnel interface on the PE router and specify the name of the routing instance:

```
[edit]
interfaces {
    interface-name {
        unit 0 {
            tunnel {
                source address;
                destination address;
                routing-instance {
                    destination routing-instance-name;
                }
                family inet {
                    address address;
                }
            }
        }
    }
}
```

To complete the GRE tunnel interface configuration, you need to configure the GRE interface at the [edit routing-instances *routing-instance-name*] hierarchy level under the appropriate routing-instance:

```
[edit]
routing-instances {
    routing-instance-name {
        interface interface-name;
    }
}
```

Configure the GRE Tunnel Interface on the CE Router

Configure the GRE tunnel interface on the CE router as follows:

```
[edit]
interfaces {
    interface-name {
        unit 0 {
            tunnel {
                source address;
                destination address;
            }
            family inet {
                address address;
            }
        }
    }
}
```

Configure an ES Tunnel Interface for Layer 3 VPNs

An ES tunnel interface allows you to configure an IP Security (IPSec) tunnel between the PE and CE routers of a Layer 3 VPN. The IPSec tunnel can include one or more hops.

To configure an ES tunnel interface between the PE and CE routers of a Layer 3 VPN, complete the procedures in the following sections:

Configure the ES Tunnel Interface on the PE Router on page 108

Configure the ES Tunnel Interface on the CE Router on page 109

Configure the ES Tunnel Interface on the PE Router

Configure the ES tunnel interface on the PE router as follows:

```
[edit]
interfaces {
    interface-name {
        unit 0 {
            tunnel {
                source address;
                destination address;
            }
            family inet {
                address address;
                ipsec-sa security-association-name;
            }
        }
    }
}
```

By default, the tunnel destination address is assumed to be in the default Internet routing table, inet.0. For IPSec tunnels using manual security association (SA), if the tunnel destination address is not in the default inet.0 routing table, you need to specify which routing table to search for the tunnel destination address by configuring the routing-instance statement. This is the case if the tunnel encapsulating interface is also configured under the routing instance.

```
[edit]
interfaces {
    interface-name {
        unit 0 {
            tunnel {
                source address;
                destination address;
                routing-instance {
                    destination routing-instance-name;
                }
                family inet {
                    address address;
                    ipsec-sa security-association-name;
                }
                family mpls;
            }
        }
    }
}
```

For IPSec tunnels using dynamic SA, the tunnel destination address must be in the default Internet routing table, inet.0.

You also need to configure the ES interface at the [edit routing-instances *routing-instance-name*] hierarchy level for the appropriate routing instance:

```
[edit]
routing-instances {
    routing-instance-name {
        interface interface-name;
    }
}
```

Configure the ES Tunnel Interface on the CE Router

Configure the ES tunnel interface on the CE router as follows:

```
[edit]
interfaces {
    interface-name {
        unit 0 {
            tunnel {
                source address;
                destination address;
            }
            family inet {
                address address;
                ipsec-sa security-association-name;
            }
        }
    }
}
```

For more information about how to configure tunnel interfaces, see the *JUNOS Internet Software Configur ation Guide: Interf aces and Class of Service* .

For more information about how to configure IPSec interfaces, see the *JUNOS Internet Software Configur ation Guide: Getting St arted*.

Configure IPSec between PE Routers Instead of MPLS

A conventional Layer 3 BGP/MPLS VPN requires the configuration of MPLS LSPs between the PE routers. When a PE router receives a packet from a CE router, it performs a lookup in a specific VRF table for the IP destination address and obtains a corresponding MPLS label stack. The label stack is used to forward the packet to the egress PE router, where the bottom label is removed and the packet is forwarded to the specified CE router.

You can provide Layer 3 BGP/MPLS VPN service without an MPLS backbone. Instead of configuring MPLS LSPs between the PE routers, you configure GRE and IPSec tunnels between the PE routers. The MPLS information for the VPN (the VPN label) is encapsulated within an IP header and an IPSec header. The source address of the IP header is the address of the ingress PE router. The destination address has the BGP next hop, the address of the egress PE router.


Note

The IPSec tunnel requires the use of an ES PIC. The GRE tunnel requires the use of a Tunnel Services PIC.

To configure IPSec between PE routers, complete the following:

1. Configure an IPSec tunnel between the PE routers. The source address is that of the ingress PE router, and the destination address is that of the egress PE router:

```
[edit interfaces]
es-interface-name {
    unit unit-number {
        tunnel {
            source source-address;
            destination destination-address;
        }
        family inet {
            ipsec-sa sa-esp-dynamic;
            address address;
        }
        family mpls;
    }
}
```

2. Configure IPSec on the PE router. For information about how to configure IPSec, see the *JUNOS Internet Software Configuration Guide: Getting Started*.

3. Configure a GRE tunnel between the PE routers. Again, the source address is that of the ingress PE router, and the destination address is that of the egress PE router:

```
[edit interfaces]
gr-interface-name {
    unit unit-number {
        family inet {
            address address;
        }
        family mpls;
        tunnel {
            source source-address;
            destination destination-address;
        }
    }
}
```

4. Configure BGP between the PE routers:

```
[edit protocols]
bgp {
    group pe {
        type internal;
        local-address local-address;
        family inet {
            unicast;
        }
        family inet-vpn {
            unicast;
        }
        peer-as as-number;
        neighbor address;
    }
}
```

5. Configure the routing instance:

```
[edit]
routing-instances {
    routing-instance-name {
        instance-type vrf;
        interface interface-name;
        route-distinguisher address;
        vrf-import import-policy-name;
        vrf-export export-policy-name;
        protocols {
            bgp {
                group routing-instance-name {
                    type external;
                    peer-as as-number;
                    as-override;
                    neighbor address;
                }
            }
        }
    }
}
```

6. Configure the policy options:

```
[edit]
policy-options {
    policy-statement import-policy-name {
        term 1 {
            from {
                protocol bgp;
                community community-name;
            }
            then accept;
        }
        term 2 {
            then reject;
        }
    }
    policy-statement export-policy-name {
        term 1 {
            from protocol [ bgp direct ];
            then {
                community add community-name;
                accept;
            }
        }
        term 2 {
            then reject;
        }
    }
community community-name members target:target;
}
```

Configure Packet Forwarding for VPNs

The procedure for configuring packet forwarding in Layer 3 VPNs is identical to the one described for Layer 2 VPNs in “Configure Packet Forwarding for VPNs” on page 26.